PATTERNS OF POWER · A FRAMEWORK PAPER

LINE 12B

Why exposure does not stop a war — the machinery by which a self-governing state disciplines its own conduct abroad, the three sensors it has removed since 1968, and the finding that what matters is not how fast the truth arrives but whose ledger it lands on

A paper in the Patterns of Power series · Thomas J. Bryan with Claude (Anthropic) · held as tested, not revealed

The claim under examination, and the discipline of this paper

There is a belief, widely held and rarely stated as a proposition, that the decisive check on a democracy's conduct abroad is disclosure. On this account the machinery runs: an investigative reporter obtains what was hidden; the public learns it; the public is appalled; and the conduct stops or is punished. The reporter is the load-bearing part. The obstacle is secrecy. The remedy is light.

The belief is not foolish, it is held by serious people, and it has real victories behind it. This paper takes it seriously enough to test it, and finds that it describes a mechanism which exists, works occasionally, and is not the mechanism that actually governs. The parts that do the governing are elsewhere, and most of them are not informational at all.

Two disciplines are observed throughout, and stating them here is not throat-clearing but a condition of the analysis being any use.

This paper takes no position on whether any particular military operation is justified. It is about machinery — how a constraint transmits or fails to transmit, and why — and it should be equally usable by a reader who supports a given intervention and one who opposes it. A theory of restraint that only works if you already agree about the war is not a theory; it is an argument for the war, wearing a lab coat.

And it assumes no prior reading. Every term is defined where it appears, every scholar is introduced with what they argued and why it bears, and every figure that carries weight has been checked.

A minimum vocabulary

Four terms, defined once.

The seam. Between what a government says about itself and what it actually does there is always a gap. The seam is that join. Every durable order does continuous work to keep the two sides of it from being held in view, at the same moment, by the same people.

The cost column. Every arrangement imposes costs on somebody. The cost column is the register of those costs *as the deciding party experiences them*. Costs that are entirely real but do not appear in that register — because they fall on people outside it — do not enter the decision. They are not concealed from the accounting; they are in a currency the accounting does not accept.

The sensor. A democracy's mechanism for pricing its own conduct is the vote, and the vote registers what the voter bears. Any institution that transfers a distant cost onto a domestic constituency — conscription, taxation, casualties among one's own, liability in one's own courts — is a *sensor*: a device that puts a foreign cost where the domestic voting mechanism can read it. Remove the sensors and the vote goes on functioning perfectly, measuring nothing.

Incidence. Borrowed from public finance, where it names the question of who actually bears a tax as distinct from who nominally pays it. Here it names the corresponding question about a disclosure: not *who learns of a cost* but *on whom the correcting of it would fall*. This paper's finding is that incidence, not information, is the governing variable.

The press does not set the range of the argument

Begin with what political scientists have actually found, because the folk theory has been tested for forty years and has not done well.

The central result is the **indexing hypothesis**, put forward by W. Lance Bennett in 1990. Its claim is that mainstream news calibrates — indexes — the range of opinion it presents to the range of opinion already being voiced by officials with institutional power. This is not a theory of cowardice or corruption. It follows from ordinary professional norms about who counts as an authoritative source and what counts as a live controversy: a reporter covering a policy fight covers the sides that exist in the institutions that decide it. The consequence is severe. When the governing class is united on a question, coverage stays inside that unity no matter what any individual reporter uncovers, because there is no institutional voice to carry the counter-position and the story has nowhere to attach. When the governing class divides, coverage opens. Reporting is therefore, in the main, a *dependent* variable. The hypothesis has been productive rather than definitively proved — Bennett's original formulation has been

criticised as difficult to falsify, and studies since have found indexed coverage in some cases and not others — and it should be held as a strong regularity rather than a law.

The finding was established against the hardest possible case. Daniel Hallin's *The Uncensored War* (1986) examined Vietnam — the war that everyone, including many of the journalists involved, believed the press had ended. Working through the actual coverage, Hallin found that critical reporting **followed** division among American elites rather than producing it. Hallin's model divides political discourse into three regions. In the **sphere of consensus**, propositions are treated as settled and not argued at all; journalistic balance does not apply because there is nothing to balance. In the **sphere of legitimate controversy**, positions are treated as genuinely contested and the norms of balance operate. In the **sphere of deviance**, positions are treated as outside serious consideration and are not aired even to be refuted. Journalism does not decide which sphere a war occupies. The behaviour of political elites does — and a war that sits in the sphere of consensus will be reported factually, accurately, in detail, and without consequence.

The optimistic counter-theory was tested too. After the interventions of the 1990s a body of work grew up around the so-called **CNN effect** — the proposition that vivid coverage of distant suffering compels governments to act. Piers Robinson examined it directly in *The CNN Effect* (2002) and found the influence weak and heavily conditional: coverage moves policy mainly where policy is *already uncertain* and elite opinion *already divided*. Which is to say the reporting is effective when the door is open, and the door being open is the variable to explain.

Set that against the folk theory and the difference is precise. The folk theory holds that exposure creates the political condition for constraint. The evidence says exposure operates *within* a political condition it does not create.

A controlled comparison: two atrocity photographs

The cleanest way to isolate the operative variable is to hold the instrument constant and vary everything else. Photography of atrocity gives an unusually good natural experiment, because the same instrument was used twice, a century apart, by campaigners of comparable seriousness, with opposite results.

Congo, 1904. From 1885 the Congo basin was the personal possession of Leopold II, King of the Belgians — not a Belgian colony but the private domain of a European

monarch, held under an international instrument, the Berlin Act, whose signatories had guaranteed free trade throughout the basin. Leopold's administration extracted wild rubber through a quota system enforced by hostage-taking, mutilation and killing on a scale whose demographic consequences remain contested and are by any estimate enormous. The campaign against it began with a shipping clerk: E. D. Morel, employed by the Liverpool firm that handled the Congo traffic, noticed from the manifests that ships returned laden with rubber and ivory and sailed out carrying not trade goods but weapons and officials — from which it followed that nothing was being *bought* in the Congo. The British consul Roger Casement produced an eyewitness report in 1904. Morel founded the Congo Reform Association with the Liverpool trader John Holt, and the campaign made pioneering use of the photograph of mutilated Congolese. It worked. Leopold surrendered the Free State to the Belgian parliament in 1908; the association wound up in 1913.

Abu Ghraib, 2004. Photographs taken by American soldiers of the abuse of Iraqi detainees at a prison outside Baghdad became public in April 2004 and circulated globally within days — a far faster, wider and more visceral dissemination than anything available to Morel. The consequences: eleven enlisted soldiers convicted by court-martial, the highest-ranking a staff sergeant, with sentences up to ten years. One officer, Lieutenant Colonel Steven Jordan, was court-martialled; he was acquitted of every charge relating to detainee mistreatment, convicted only of disobeying an order not to discuss the investigation, and that single conviction was subsequently set aside, leaving him with an administrative reprimand. The commanding officer of the prison, Brigadier General Janis Karpinski, was demoted administratively. No senior officer and no civilian official faced criminal liability. The administration was returned to office six months later. The one substantial civil verdict — forty-two million dollars awarded to three Iraqi former detainees against a private interrogation contractor — arrived in November 2024, more than twenty years after the conduct, and against a company rather than the government.

Same instrument. Opposite result. Now vary the one thing that differs.

In the Congo case the target was a **foreign sovereign**; no British voter had to change anything. A **concentrated domestic commercial interest** — traders excluded from the basin by Leopold's monopoly in breach of the Berlin Act — wanted the correction and financed the campaign. The load-bearing public argument was **free trade**. The cost of the reform fell on a Belgian king.

In the Abu Ghraib case the target was **one's own government and army**. No domestic faction's interest was served by a reckoning. The correction, if made, would have landed on the people making it.

Which yields the rule, and it is the rule this paper exists to state:

> **The exposure threshold is a function of whose ledger the correction lands on.**

A state audits a rival's periphery in years and its own across decades. It is the same sensor in both cases. It simply reads only what its own constituents pay.

The pattern holds in the other famous instance. Daniel Ellsberg leaked the Pentagon Papers to the press in June 1971 expecting the disclosure to end the war; the war continued for years afterward. And note what the Papers actually contained: a documentary history of *previous* administrations' deceptions. The disclosure indicted men substantially out of power. It was, in the terms of this paper, an exposure whose incidence fell on nobody currently deciding anything.

The three sensors, and when they were removed

Here the analysis leaves the media entirely, because the largest changes in the last sixty years are not informational. Three institutions that once transferred the cost of foreign war onto the American domestic ledger have been removed, and none of them was removed for reasons that had anything to do with foreign policy.

Conscription. The draft ended in 1973 and the United States moved to an all-volunteer force. Whatever else conscription is, it is the most powerful cost-transfer mechanism a state possesses: it places the risk of a distant war inside a very large number of households that would otherwise never encounter it, and it does so involuntarily, which is what makes it register. It is the principal reason Vietnam moved American opinion in a way that nothing since has approached. The sociologist Charles Moskos put the corollary sharply: once the children of the governing class no longer had to serve, tolerance for casualties among those who did should be expected to fall, because the people setting the policy no longer stood in the line of fire.

War finance. Every major American war from the Civil War through Korea was paid for substantially by taxation, and the taxes were visible. Korea saw increases on the order of four per cent of GDP. Vietnam's tax moment came in 1968, and it produced the

artefact this paper is named for. The Revenue and Expenditure Control Act of 1968 imposed a temporary ten per cent surcharge on individual and corporate income tax — effective for individuals from 1 April 1968, exempting the lowest incomes, raising revenue on the order of one per cent of GDP, and contributing to the last federal budget surplus until 1998. Its administrative form is the thing worth holding. The 1968 Form 1040 acquired a new line. A taxpayer computed his ordinary liability on line 12a, consulted a surcharge table, and entered the addition on **line 12b**.

That is what a cost transfer looks like as a physical object: a war, on a tax return, on a line, once a year, for every filer in the country.

The post-2001 wars were financed differently, and the difference is not a matter of degree. They were paid for essentially entirely by borrowing — the first major American wars since the Revolution to be conducted with no wartime tax increase and no offsetting reduction in other spending. Taxes were *cut* in 2001 and again in 2003, and again in 2017; federal receipts fell from roughly 18.8 per cent of GDP in 2001 to about 16.2 per cent by 2020 while the wars ran. The spending itself was routed through a separate supplemental budget rather than integrated into the regular defence budget, which removed the internal trade-offs that would otherwise have forced the Pentagon and the appropriators to choose between the war and something else. Sarah Kreps's work on war taxation and Rosella Cappella Zielinski's on war finance both reach the same structural conclusion: shielding the public from the direct cost of a war weakens the institutional constraints that historically kept democratic wars short. There is no line 12b.

Casualties among one's own. The most robust finding in the literature on war and public opinion is John Mueller's. Examining polling across Korea and Vietnam — first in a 1971 article and then in *War, Presidents and Public Opinion* (1973) — Mueller found support for both wars beginning high and declining as a *logarithmic* function of cumulative American dead, with a curve remarkably similar across the two conflicts. The logarithm matters: it means each additional death moves opinion less than the one before, so early losses are disproportionately powerful. The finding has been refined rather than overturned — Scott Gartner and Gary Segura showed in 1998 that *marginal* casualties, the rate at which deaths are accumulating at a given moment, add explanatory power that cumulative totals alone miss — but the core holds. American deaths are the strongest known predictor of declining American support for an American war.

A campaign conducted from standoff range, by aircraft, missiles and remotely piloted systems, drives that number toward zero. This is not framed here as a criticism of force protection, which is a legitimate and indeed obligatory military objective. It is a structural observation: the single most reliable domestic sensor for the cost of a foreign war is, as an incidental consequence of protecting one's own troops, engineered out of the system.

The drift, and why it is not a conspiracy

It matters a great deal to state the next point precisely, because the material invites a conspiratorial reading it does not support.

None of the three sensors was removed in order to insulate foreign policy from domestic constraint. Conscription ended because it was unpopular, inequitably administered, and thought — with reason — to produce a worse army than volunteering. The 1968 surcharge lapsed because it was legislated as temporary and because a permanent war tax is a thing no legislature enacts if it can avoid one. Standoff warfare was developed to win engagements and to keep soldiers alive. Each removal had its own justification, most of them good, and each was decided in its own domain by people who were not thinking about the aggregate at all.

The result is nonetheless an aggregate: a state that retains its full constitutional machinery for pricing its own conduct and has, without anyone deciding to, disconnected the instruments that fed the machinery its inputs. This is drift, not design. It is also the harder problem, because there is no decision to reverse and nobody to hold responsible — which is precisely why it persists, and why an analysis that goes looking for a culprit will find none and conclude, wrongly, that there is nothing there.

The clause that exempts, translated into law

Colonial orders held their coercion and their professed principles together by a clause: the governed were not excluded from the moral community, they were held to be *not yet ready* for its full application — in nonage, at an earlier stage, not capable yet of self-government. Its power lay in three properties. It could be held sincerely by a decent person. It was shaped like an empirical claim, so it invited evidence and felt like a finding

rather than a prejudice. And the assessment of readiness was performed by the party that benefited from a negative answer, so in practice the answer was always *not yet*.

The modern equivalent is not developmental. It is juridical, and it does the same work with the same three properties.

Consider the American legal architecture for targeted killing, as set out in the Justice Department white paper that became public in 2013. On the question of when a threat is sufficiently imminent to justify lethal force, the paper states:

> The condition that an operational leader present an "imminent" threat of violent attack against the United States does not require the United States to have clear evidence that a specific attack on U.S. persons and interests will take place in the immediate future.

Instead, the determination is committed to an informed, high-level official, and the threat posed by the relevant organisation is characterised as *continual*, on the reasoning that an adversary which does not mass forces or wear uniforms cannot be caught at a recognisable end-stage of planning, so that waiting for one would be to accept an unacceptable risk.

Read structurally rather than polemically. That reasoning is coherent; there is a real problem it is responding to, and lawyers of considerable seriousness have defended it. But note the three properties. It is **sincerely holdable** — the officials making these determinations are, in the main, careful professionals who believe what they are asserting. It is **shaped like an empirical finding** — the question becomes whether the assessment was accurate, not whether the killing was permissible. And the assessment is **performed by the party that benefits from the answer**, with the underlying intelligence classified, so that the empirical question it has become is one that only the deciding party is equipped to argue.

That last property is the whole function. The clause converts a moral question, on which any citizen has standing to hold a view, into a classified technical question, on which almost none do. And a citizen who has been told that the matter turns on an intelligence assessment he cannot see has been told, accurately and without any deception, that he has nothing to contribute.

The same move appears in a lower register whenever a legal threshold has to be met. The 2019 presidential veto message on Yemen turns on the assertion that, apart from

counterterrorism operations, the United States "is not engaged in hostilities in or affecting" Yemen — a definitional claim about the word *hostilities* that decides the statutory question before it is reached. Whether that reading was right is a legitimate argument. What matters here is the form: the exemption does not deny the conduct or the principle. It relocates the dispute into a definition.

The classification map is a threshold map

An earlier paper in this series established a rule from the colonial archive that transfers exactly. In Northern Nigeria in 1902–03, the British High Commissioner Frederick Lugard concealed the preparations for a war from the Colonial Office in London — refusing to use the telegraph, sending despatches by sea that would arrive too late to matter — because the Colonial Office could have stopped him. The same man published, in an annual report laid before Parliament and sold for one shilling and threepence, an account of conduct that flatly contradicted Britain's professed anti-slavery commitments, because Parliament would not act on it. He was right on both counts. The rule that falls out is: **what a regime bothers to hide is a map of what its audience would actually punish.**

The American pattern has the same shape. Consider what is and is not withheld.

Not withheld: the fact of strikes, which are announced. Not withheld: notification to Congress, which the War Powers Resolution of 1973 requires within forty-eight hours of introducing forces into hostilities, and of which the War Powers Resolution Reporting Project counted 135 instances as of July 2026. Not withheld: broad legal position papers, eventually.

Withheld: the operative legal memoranda, the targeting criteria, the underlying intelligence, the casualty assessments. And withheld, crucially, from the institution that could act on them.

The judicial branch is the clearest case, because the mechanism is visible on the record. In 2010 the father of Anwar al-Aulaqi, an American citizen reportedly placed on a targeting list, sought an injunction in federal court. In *Al-Aulaqi v. Obama* the district court dismissed the suit on 7 December 2010 without reaching the merits, holding that the father lacked standing to assert his adult son's rights and that the questions raised were political questions committed to the elected branches; the government had

additionally invoked the state secrets privilege, arguing that the case could not be litigated without disclosing classified information. Standing, political question, state secrets: three doctrines, none of them invented for this purpose, each perfectly respectable in its own domain, and their combined effect is that the body with power to compel disclosure is the body that receives none.

That is Lugard's asymmetry, reproduced in the American constitutional order without anyone having designed it that way: conceal from the institution that could stop you, publish to the institutions that will not.

The votes are the measurement

The most useful instrument for reading a legislature's actual threshold is not what its members say but the gap between what they will vote for when the vote binds and when it does not. That gap is directly observable.

Yemen, 2019. A joint resolution directing the removal of United States forces from hostilities in Yemen — a *binding* instrument, which if enacted would have had the force of law — passed the Senate 54–46 on 13 March and the House 247–175 on 4 April, with sixteen Republicans joining the majority. It was vetoed on 16 April, and the Senate failed to reach the two-thirds needed to override on 2 May. On the surface this reads as a defeat for the mechanism. But note what did work in the same period: a condition attached to the previous National Defense Authorization Act restricting in-flight refuelling of coalition aircraft appears to have produced an actual change in American conduct. The symbolic vote failed. The statutory condition, embedded in a bill that had to pass, bit.

Iran, July 2026. On 23 July the House approved, 214–208, a *concurrent* resolution calling on the President to end American involvement in the conflict — an instrument which, as concurrent resolutions do, lacks the force of law. On the same day the Senate fell short, 49–47, of even discharging from committee a *binding* joint resolution to the same end.

Hold those two numbers together. A chamber-majority for the version that does nothing; short of a majority for the version that would do something. The gap between the binding and the nonbinding tally is a direct numeric reading of a legislature's threshold — the price at which its members will act rather than merely position.

And there is a reason the binding vote is harder, which returns the argument to its centre. **A recorded binding vote is itself a cost transfer.** It converts a distant foreign cost into a personal domestic liability: a specific member's name, on a specific record, attached to a specific outcome, retrievable by a challenger. That is exactly why it is difficult to obtain — and exactly why it is the thing worth fighting for. The instrument that constrains and the instrument that is resisted are the same instrument, for the same reason.

What exposure actually yields

The strongest modern test of the folk theory is also its finest recent achievement, and it deserves to be described accurately rather than used.

Between 2017 and 2021 the reporter Azmat Khan pursued the record of American air operations in Iraq, Syria and Afghanistan through Freedom of Information Act requests and, when those failed, through litigation against the Department of Defense. She obtained more than 1,300 formerly secret military records — including internal civilian-casualty assessments — and paired them with ground investigation at more than a hundred strike sites, comparing what the assessments claimed against what could be established from survivors, hospital records, death certificates and physical evidence. The resulting series, published in *The New York Times* in December 2021 as "The Civilian Casualty Files," established a substantial and systematic gap between official accounting and what had happened, and won the 2022 Pulitzer Prize for International Reporting. It is, on any measure, first-rate investigative work performed at real personal and institutional cost.

Now trace its output. Khan's earlier reporting contributed to Section 936 of the National Defense Authorization Act for Fiscal Year 2019, the first congressionally mandated standards for civilian-casualty policy. The *Civilian Casualty Files* prompted proposed legislation — a Protection of Civilians in Military Operations Act, a Civilian Harm Transparency Act — a Senate Judiciary Committee hearing on drone strikes, letters from members, a documented shift in sentiment among defence officials, and a body of subsequent scholarship and advocacy.

That is a real record and this paper does not diminish it. But read the *shape* of the yield. The exposure produced **standards, transparency requirements, hearings, and**

reporting obligations. It did not produce a cessation, a prosecution, or a change in the authority under which the strikes were conducted.

Which suggests a structural observation that generalises across the modern cases:

> **The characteristic yield of a successful modern exposure is a reporting requirement — an instrument whose output is more exposure.**

This is not cynicism about the reformers, who wanted more and took what was available. It is an observation about incidence. A reporting requirement is the cheapest concession an institution can make, because its cost is administrative rather than operational: it obliges the institution to produce documents, and producing documents is a thing institutions are extremely good at. It transfers no cost to any domestic constituency. It creates no liability. It restrains no action. And it closes the loop back onto the very mechanism whose insufficiency produced the demand — the audit is improved, and the audit was never the binding constraint.

The revision this case forces

The companion paper to this one examined why the atrocities of European colonialism required almost no censorship, and concluded that the liberal order's fault was **latency**: that its audit machinery is real but runs at roughly a century's lag, long enough for every beneficiary to die comfortable. It prescribed shortening the lag.

The American case shows that diagnosis to be partly wrong, and the correction is the most important thing in this paper.

Latency has already collapsed. Abu Ghraib was global within days. The *Civilian Casualty Files* arrived while the campaigns they described were ongoing. Strike reporting operates at something close to real time, drawing on satellite imagery, open-source verification and immediate testimony that no nineteenth-century reformer could have imagined. By the standard of any colonial case in the corpus, the modern audit is astonishingly fast.

And constraint has not increased in proportion. It may not have increased at all.

So:

> **Latency reduction is necessary and not sufficient. The operative variable is not lag but incidence.** A fast audit whose correction lands on nobody produces the same outcome as a slow one — it simply produces it sooner and more legibly.

The century-lag finding was measuring a symptom. What made the colonial audit slow and the modern audit fast is the technology of disclosure; what made both of them non-binding is the same in each era, and it is the incidence of the correction.

The strongest counter-position

Three objections, each serious, stated at strength before being answered.

The first is that this analysis is demoralising and therefore harmful. People do the hardest and most dangerous work in journalism on the belief that exposure matters; an argument that it does not is, whatever its analytic merits, a contribution to the very inattention it describes. If the reporters had believed this they would not have gone.

The second is the long tail. Exposure may not stop a war, but it constrains at the margin — rules of engagement tightened, targets not struck, careers ended quietly, the knowledge among officers that a thing may become public. It also builds the documentary record from which later law is made; the British abolition of the slave trade rested on decades of pamphlets, testimony and accumulated public knowledge before any vote was won, and to judge that literature by whether it stopped a slave ship in the year of its publication is to apply the wrong time horizon.

The third is that the counterfactual is unobservable. One cannot count the strikes not taken, the operations not authorised, the policies abandoned in draft because someone in the room said *this will end up in the newspaper*. A study that measures only what happened will systematically undervalue a deterrent.

All three are granted, and the third in particular cannot be dismissed. But none of them touches the claim actually made here, because the claim is not that exposure is futile. It is that exposure is **insufficient**, and specifically that it is the *raw material* of constraint rather than the *instrument* of it.

That distinction survives all three objections and is not demoralising but operational. It tells the person doing the difficult work where the leverage sits. Khan's 1,300 records are precisely what a liability statute, a standing provision, a discharge petition or a forced recorded vote would have to run on; without the records, none of those instruments has

anything to bite. The argument is not *stop reporting*. It is: **the reporting is necessary and the instrument is separate, and the years spent waiting for the reporting alone to work are years in which the instrument was not being built.**

Reading the case against the framework

Five contributions follow.

One. Incidence supersedes latency. The governing variable in whether a disclosure constrains conduct is not how quickly the truth arrives but on whose ledger the correction would fall. Where the correction serves a concentrated domestic interest or costs the deciding public nothing, exposure moves policy quickly; where it would land on the deciding public itself, exposure produces acknowledgement without action, regardless of speed, vividness or volume. This supersedes the latency finding of the companion paper, which was measuring a symptom of the same cause.

Two. Sensor removal is drift, not design. The instruments that transfer foreign costs onto a domestic ledger — conscription, war taxation, exposure of one's own to risk — are removed one at a time, in separate domains, for reasons internal to each, most of them good. No one decides to insulate foreign policy from domestic constraint. The insulation is the aggregate of decisions taken about other things, which is why it is stable: there is no decision to reverse and no author to hold responsible. Any analysis that goes looking for an author will find none and wrongly conclude there is nothing there.

Three. The deferral clause survives translation. The colonial exemption — *not the same rules, not yet* — reappears in mature legal orders as a technical vocabulary: definitional thresholds, standards of imminence, contested categories of hostilities. The three properties that made the colonial version durable are preserved intact: sincerity, empirical form, and assessment by the interested party. The translation adds a fourth and stronger property, which is that the empirical question the moral question has become is classified, so that the only parties competent to argue it are the parties that decided it.

Four. Exposure's characteristic yield is procedure, and procedure is the cheapest available concession. A successful modern exposure typically produces standards, hearings, transparency requirements and reporting obligations, because those cost an institution administrative effort and nothing else — no liability, no operational restraint, no transfer to any constituency. They also close the loop back onto disclosure,

improving the audit that was never the binding constraint. A reform that yields only procedure is a measurement of the threshold, not a crossing of it.

Five. The binding/nonbinding vote gap is a direct instrument. Where a legislature can vote on the same question in symbolic and enforceable form, the difference between the two tallies measures its members' actual threshold with unusual precision, because the only variable is whether the vote creates personal liability. And the reason such votes are resisted is the reason they work: **a recorded binding vote is itself a cost transfer**, converting a distant foreign cost into a retrievable domestic one attached to a name.

And the design instruction that follows, which is the practical point of the paper. **Do not aim at the audience. Aim at the incidence.** Liability with a domestic forum. Standing that survives dismissal. Jurisdiction. Conditions embedded in bills that must pass rather than resolutions that need not. Forced recorded votes under privileged procedures. Disclosure with a price attached rather than disclosure alone. A war surtax with a line on the return, which is the oldest and crudest of these and possibly still the best, because it is the only one that reaches every voter once a year without requiring any of them to read anything.

Not persuasion. Accounting.

What would refute this paper

Two demonstrations would break it.

The first. A case in which disclosure alone halted a state's ongoing and materially beneficial conduct abroad — at the point of exposure, on the strength of the exposure, with no cost transferred to any domestic constituency, no elite faction's interest served by stopping, and no rival power embarrassed. Not a cessation that arrived after the conduct had ceased to pay. Not a reform some faction wanted anyway and used the disclosure to obtain. Not a withdrawal already priced by someone with a budget. One clean case breaks the incidence rule.

The second. A demonstration that the three sensors are not load-bearing: a sustained, costly foreign campaign that generated durable domestic constraint under conditions of an all-volunteer force, deficit financing, and negligible casualties among the acting state's own personnel. If constraint can be produced with all three sensors

disconnected, then something other than incidence is doing the work and the mechanism proposed here is wrong.

Both are findable in principle. Until one is found, the reading stands: the machinery is intact, the audit is faster than it has ever been, and the fault is not in what the public is permitted to know but in what any of them are made to pay.

A note on sources. Quotations from United States government documents are given exactly and are not subject to copyright: the passage on imminence is from the Department of Justice white paper on the lawfulness of lethal operations against senior operational leaders of al-Qa'ida or associated forces, which became public in February 2013; the characterisation of American involvement in Yemen is from the presidential veto message accompanying S.J.Res.7, 16 April 2019.

Modern copyrighted scholarship is paraphrased and credited, never quoted at length. The indexing hypothesis is W. Lance Bennett's, "Toward a Theory of Press-State Relations in the United States," *Journal of Communication* 40 (1990); the objection that its original formulation is difficult to falsify, and the mixed findings of subsequent tests, are reported because they are real. The three spheres and the Vietnam finding are from Daniel Hallin, *The Uncensored War: The Media and Vietnam* (1986). The conditional character of the CNN effect is from Piers Robinson, *The CNN Effect: The Myth of News, Foreign Policy and Intervention* (2002). The casualties hypothesis is John Mueller's, from "Trends in Popular Support for the Wars in Korea and Vietnam," *American Political Science Review* 65 (1971), and *War, Presidents and Public Opinion* (1973); the marginal-casualties refinement is Scott Gartner and Gary Segura, "War, Casualties, and Public Opinion," *Journal of Conflict Resolution* 42 (1998), and it is a genuine qualification rather than a footnote. The war-finance argument follows Sarah Kreps's work on war taxation and Rosella Cappella Zielinski's on war financing, together with the Costs of War project at Brown University and the "ghost budget" analysis of separate-account war funding; Charles Moskos's observation about elite exposure to casualties is credited to him.

Figures. The 1968 surcharge — ten per cent, Revenue and Expenditure Control Act of 1968, effective 1 April 1968 for individuals and 1 January 1968 for corporations, with low incomes exempt, raising roughly one per cent of GDP — is given as the tax-history literature gives it, as is the detail of line 12b on the 1968 Form 1040. The federal-receipts figures (18.8 per cent of GDP in 2001 falling to 16.2 per cent by 2020) are as reported in the ghost-budget literature and are approximate. Aggregate costs of the post-2001 wars are deliberately omitted, because the published estimates differ by trillions depending on whether veterans' care and compounding interest are included, and no single figure would be a finding. Congo mortality figures are omitted for the same reason.

Vote tallies are given as reported: S.J.Res.7 passed the Senate 54–46 on 13 March 2019 and the House 247–175 with one present on 4 April 2019, was vetoed on 16 April, and the Senate override attempt failed on 2 May. The votes of 23 July 2026 — a House concurrent resolution approved 214–208, and a Senate motion to discharge a binding joint resolution failing 49–47 — are reported from contemporaneous coverage on the day after; readers should verify against the official record, since these are days old at the time of writing and the situation they concern is ongoing.

The Abu Ghraib outcomes — eleven enlisted convictions with the highest rank a staff sergeant; the court-martial, acquittal on all abuse charges, single collateral conviction and subsequent setting-aside in the case of Lieutenant Colonel Steven Jordan; the administrative demotion of Brigadier General Janis Karpinski; and the $42 million civil verdict against a private contractor in November 2024 — are as reported in the contemporaneous and legal-historical record. *Al-Aulaqi v. Obama*, 727 F. Supp. 2d 1 (D.D.C. 2010), was dismissed on 7 December 2010 on standing and political-question grounds without reaching the merits, with the state secrets privilege also asserted. The count of 135 War Powers Resolution reports is from the War Powers Resolution Reporting Project as of July 2026 and will have changed.

Azmat Khan's "The Civilian Casualty Files" (*The New York Times*, December 2021; 2022 Pulitzer Prize for International Reporting) is described from the public record of the project and its documented effects, including Section 936 of the FY2019 National Defense Authorization Act, subsequent proposed legislation, and a Senate Judiciary Committee hearing; the reporting itself is paraphrased and not quoted. The Congo Reform Association material follows the campaign's standard history, with the commercial and free-trade motivations alongside the humanitarian ones stated as the current reading of the scholarship rather than a settled one.

The framework vocabulary — the seam, the cost column, the sensor, incidence, and the rule that what a regime conceals maps what its audience would punish — is Thomas J. Bryan's, developed across the companion papers of this collection and defined in the text here so that this paper stands alone. It takes no position on the merits of any military operation described. It is a reading of a mechanism across a finite set of cases, stated so that it can be checked and, where it is wrong, broken.

Part of the collection Patterns of Power. *Tested, never revealed.*